

UMD-LM-16-011

**A Real Options Approach to Quantity and Cost
Optimization for Lifetime and Bridge Buys of Parts**

By:

Peter Sandborn and Navid Goudarzi

***Electronic System Cost Modeling Laboratory
(CALCE Electronic Products and Systems Center)***

December 2015

This research was partially sponsored by
the Naval Postgraduate School.



THIS PAGE LEFT INTENTIONALLY BLANK



Abstract:

Mismatches between part procurement life cycles (especially for electronic parts) and the life cycles of the systems that parts are used in cause systems with long manufacturing and/or support lives incur significant obsolescence management costs. While lifetime and bridge buys, which are used to manage part discontinuance (i.e., obsolescence) during the support of critical systems are simple in concept, the challenge is determining the optimum number of parts to buy and understanding their true cost. This report presents a real options approach to value lifetime and bridge buys and to determine the optimum part quantity for a lifetime or bridge buy. The approach accommodates uncertainties in demand, holding costs and end of support dates.

Acknowledgment:

This material is based in part upon work supported by the Naval Postgraduate School Acquisition Research Program under Grant No. N00244-14-1-0060 and the Office of the Secretary of Defense. The views expressed in this report do not necessarily reflect the official policies of the Naval Postgraduate School nor does mention of trade name, commercial practices, or organizations imply endorsement by the US Government.

Authors:**Navid Goudarzi**

CALCE Center for Advanced Life Cycle Engineering Research
Associate, Mechanical Engineering Department University of Maryland
College Park, MD 20742
navid1@umd.edu

Peter Sandborn

CALCE Center for Advanced Life Cycle Engineering Professor,
Department of Mechanical Engineering University of Maryland
College Park, MD 20742
sandborn@umd.edu
(301) 405-3167

THIS PAGE LEFT INTENTIONALLY BLANK

1. INTRODUCTION

Background

A significant problem facing many complex systems is technology obsolescence. Technology obsolescence is defined as the loss or impending loss of original manufacturers of items or suppliers of items or raw materials (Sandborn 2008). The type of obsolescence addressed in this report is referred to as DMSMS (Diminishing Manufacturing Sources and Material Shortages) and is caused by the unavailability of technologies or parts that are needed to manufacture or sustain a system. DMSMS means that due to the length of the system's manufacturing and support life, coupled with unforeseen support life extensions, needed parts become unavailable (or at least unavailability from their original manufacturer). While DMSMS can impact hardware, software, intellectual property, and human capital, its impact on hardware, specifically electronic parts, is the focus of most existing DMSMS management.

The DMSMS type obsolescence problem is most prevalent in "sustainment-dominated" systems (Sandborn & Myers 2008) where the cost of sustaining (maintaining) the system over its support life far exceeds the cost of manufacturing or procuring the original system. Sustainment in this report refers to three things: keeping a system operational, continuing to manufacture and install versions of the original system that satisfy the original requirements, and finally manufacturing and installing versions of the original system that satisfy new and evolving requirements. Examples of sustainment-dominated systems include airplanes, military systems, medical equipment, and power plant controls and telecommunications infrastructure. These types of systems have long enough design cycles that a significant portion of the electronic technology in them are obsolete prior to the system being fielded for the first time, after which they must be supported in the field for 20+ years. For these systems, simply replacing obsolete parts with newer parts is often not a viable solution because of high re-engineering costs and the prohibitive cost of system qualification and certification.

Many part obsolescence mitigation strategies exist for managing DMSMS obsolescence once it occurs, including (Stogdill 1999): lifetime buy (also referred to as final order or Life Of Type - LOT buy), bridge buy (also referred to as last-time buy), alternative/substitute parts, aftermarket sources, emulation, re-engineering, salvage, thermal uprating, and design refresh/redesign of the system.

Lifetime and Bridge Buys

The opportunity to make lifetime buys – making a one-time purchase of all the parts that you think you will need forever, is usually offered by manufacturers of electronic parts prior to part discontinuance (in the form of a published "last order date"). Alternatively, bridge buys mean purchasing enough parts to last until a planned design refresh point in the future where the part will be designed out. Lifetime and bridge buys play a role in nearly every part obsolescence management portfolio no matter what other reactive, pro-active or strategic management plans are being followed.

Purchasing sufficient parts to meet current and future demands is simpler in theory than in practice due to many interacting influences and the complexity of multiple concurrent buys as shown in Fig. 1. Fundamentally, the lifetime buy problem can be divided into two activities: 1) demand forecasting and 2) optimizing the buy quantities based on the demand forecasted.

Forecasted demand depends on manufacturing (sales) forecasts and sustainment expectations (spares) for fielded systems – this report does not address the demand forecasting portion of the problem. The second portion of the problem is given the demand and its uncertainties, determine the number of parts that should be purchased (buy quantity), which is the focus of this report.

In practice today, the common wisdom usually used to determine buy sizes is a best guess (forecast) of demand based on projected manufacturing needs (if manufacturing is still occurring) and spares needed (based on observed or predicted failure rates) to the planned end of support (EOS) date; then

buffer that quantity by 10%-50%. Over time, the buffers often increase due to pain experienced by engineers.¹

Quantitative approaches to the lifetime/bridge buy problem have been used - given a demand forecast one can calculate the quantities of parts necessary to minimize life-cycle cost, which, depending on how you are penalized for running short or running long on parts could be substantially different than what simple demand forecasting tells you to purchase. In general, this is an asymmetric problem where the penalties for under buying parts and overbuying parts are not the same – if they were the same, the optimum quantity to purchase would be exactly the forecasted demand. For example, the penalty for under buying parts is the cost to acquire additional parts long after they became obsolete or redesign the system to use a newer part; while the penalty for overbuying parts is paying for extra parts and paying the inventory (holding) cost for those parts for a long period of time and then losing all or some of that investment.² In general, for sustainment-dominated systems, the penalty for under buying parts is significantly larger than the penalty for overbuying parts.

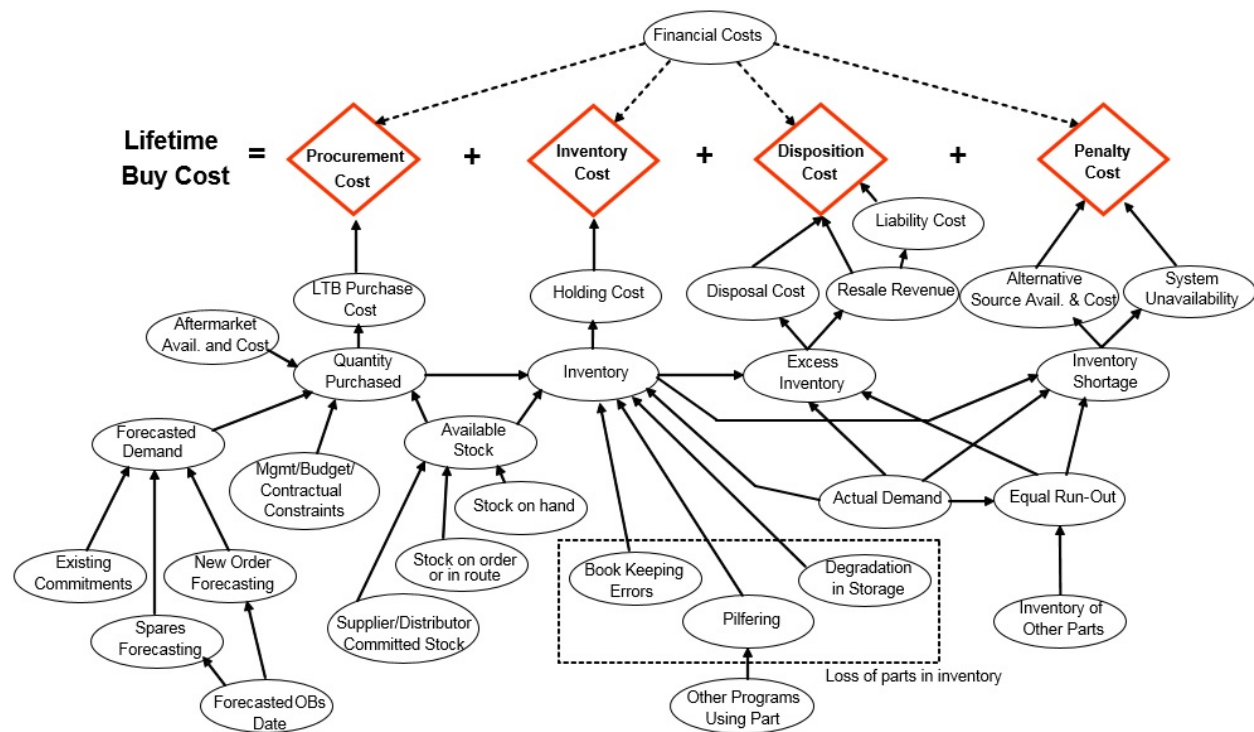


Figure 1 - Lifetime buy costs (Feng et. al. 2007).

¹ The buffers are put in place to mitigate “life extension” risk. Life extensions may take the form of: a) manufacturing the product that the part is in for longer than anticipated, b) supporting the fielded products for longer than planned, or c) design refreshes (that designed out the obsolete part) happening less frequently than planned or taking longer than planned. In most organizations, the buffers are based on “institutional knowledge” and there may be little understanding of the statistical meaning or ramifications of the lifetime/bridge buy buffer sizes that are used.

² Additionally, you may need to pay to dispose of the extra parts. The cost of disposal could be negative (if you resell the parts) or positive (paying to insure that parts are destroyed so they can't enter the counterfeit parts stream).

Existing Lifetime/Bridge Buy Analysis Approaches

In the operations research domain, lifetime buy optimization is a special case of the of the newsvendor problem.³ Extensions to the classical newsvendor problem solution exist that accommodate many different situations, but these solutions fall well short of solving real lifetime and bridge buy problems because they generally lack time dependence, i.e., they generally do not include cost of money and holding cost. In addition, a “must support” assumption is implicit in lifetime buy problems that is not generally present in simple newsvendor problems; you cannot choose not to support the system, i.e., you are not allowed to fail to fulfill the demand and therefore you must pay the penalty to purchase extra parts from a broker or redesign the system if you run out (the newsvendor is not required to do this). A discussion of the application of newsvendor problem solutions to lifetime buys appears in (Sandborn 2013).

Some treatments of the “final order” problem applicable to lifetime buy exist in the operations research literature. Existing final order models are intended for systems like complex manufacturing machinery that have long-term service contracts. To be able to provide long-term service, a manufacturer must be able to supply parts throughout the service period. The period after the machine has been taken out of production is called the *EOS* period. To avoid out-of-stock situations during the *EOS*, an initial stock of spare parts is ordered at the beginning of the *EOS*. This initial stock is called the final order. Some final order problem solutions exist (Teunter & Fortin 1999, Tunter & Haneveld 1998), but simplifying assumptions about demand profiles and fixed *EOS* dates make these solutions impractical for the treatment of DMSMS problems in real applications. The problem with existing newsvendor and operations research solutions is that they either cannot accommodate the important drivers (e.g., time), or are oversimplified to the point of uselessness for real applications. Most practical treatments of lifetime buys use discrete- event simulation (*DES*) that follows the time history of a population of parts forecasting demands, determining holding costs and associated penalties until an *EOS* date for the use of the part is reached. Such solutions can be used to determine the life-cycle cost of the buys and the optimum quantities (the quantity that minimizes the life-cycle cost). These solutions base their optimum quantity decision on the minimization of life-cycle cost.

Section 2 of this report proposes a real options approach to the life-time buy (LTB) quantity optimization problem. Section 3 provides a verification of the real options model using a stochastic *DES* solution. Both real options and stochastic *DES* accommodate uncertainties in demand, holding costs and *EOS* dates. In Section 4, the dependence of the solutions on the weighted average cost of capital (the cost of money) and uncertain *EOS* dates is discussed.

2. A REAL OPTIONS APPROACH TO LIFETIME BUY FORECASTING

This section presents a methodology to quantify a proposed option in the ROA method and to illustrate its role in obtaining the optimum LTB/bridge buy. Section 2.1 presents the valuation methodology; Section 2.2 focuses on the real options analysis of the option; and Section 2.3 provides a demonstration case.

2.1 Valuation Methodology

A real option is the right, but not the obligation, to undertake certain business initiatives, such as deferring, abandoning, expanding, staging, or contracting. For example, the opportunity to invest in an asset is a real “call” option. Real options differ from financial options in that they are not typically traded as securities, and do not usually involve decisions on an underlying asset that is traded as a financial security. Unlike conventional net present value analysis (discounted cash flow analysis) and decision tree analysis, real options offers the flexibility to alter the course of action in a real asset decision, depending on future developments.

³ The newsvendor problem seeks to find the optimal inventory level for an asset given an uncertain demand and unequal costs for overstock and understock. This problem dates back to an 1888 paper by Edgeworth (Edgeworth, 1888).

The analysis of options focuses on valuation under uncertainty. If there was no uncertainty, the value of an option would be trivial to determine. However, everything is uncertain and the future returns are generally highly asymmetric (upside \neq downside). For financial options, the important question is what should I pay to buy the option? For real options the questions are what is the value I get from the option and when do I exercise the option? Real options have been previously used in obsolescence management to assess the value of waiting to invest in new technology (Josias 2009), to assess delaying lifetime buys (Venesmaa et al. 2008), and the option to buy additional parts after the lifetime buy (Burnetas and Rtchken 2000); however, real options have never been used to assess lifetime or bridge buy sizes

In order to cast the *LTB* problem for ROA we need to define the “value” of the lifetime buy option,

$$V = CAPenalty + CALTB \quad (1)$$

where $CAPenalty$ is the penalty cost avoided (penalty cost associated with running out of parts before the demand for parts is exhausted) and $CALTB$ is the lifetime buy cost avoided (by buying and holding fewer parts). Figure 2 shows a graphical representation of (1).




Figure 2 – Simple lifetime buy (*LTB*) value formulation.

For example, if zero parts are purchased at a lifetime buy, then no penalty costs are avoided and all *LTB* costs are avoided. For a non-zero quantity, N , of lifetime buy parts purchased, there will be a greater than zero penalty cost avoided, but *LTB* costs will have to be paid. The cost avoidances in (1) are uncertain due to uncertainties in everything: the demand for parts for manufacturing and for sparing, the timing of demands, and the penalties. In order to formulate this as a real options problem we must cast the formulation in terms of time (not part count); hence, the cost avoidances and value that appears in (1) become functions of “buy-to time (T_{bt})” and the objective will be to determine the optimum point in time to buy parts for, which is the optimum time to exercise the option. Knowing the optimum exercise time, the *equivalent optimum number of parts* can be determined.

For an arbitrary buy-to time (T_{bt}) between zero to the *EOS*, (shown in Fig. 3), (1) is defined as:

$$V_i(T_{bt}) = \left[\sum_{t=0}^{EOS} (L_i(t) + P_i(t)) - \sum_{t=T_{bt}}^{EOS} (L_i(t) + P_i(t)) \right] + [LTB_i(EOS) - LTB_i(T_{bt})] \quad (2)$$

$$LTB_i(T_{bt}) = IB_i(T_{bt}) + HC_i(T_{bt})$$

where i refers to i^{th} simulated demand path, L and P refer to lump and pay per part penalties imposed after the *LTB* runs out, IB is the initial buy cost, and HC is the sum of all the holding costs for the initial

buy of parts as they are held and used up to the buy-to time. The lump penalty is a one-time charge at the *LTB* run out date, and the pay per part penalty is imposed for every part that must be procured after the *LTB* runs out. The first term in brackets in (2) is the penalty cost avoided (it includes under buy penalties due to buying less than the required *LTB* size) and the second term in brackets is the lifetime buy cost avoided (it includes the purchase of the initial buy and the holding cost for the parts up to the points where they are used).




Figure 3 - Valuation at an arbitrary buy-to time, T_{bt} .

2.2 Real Options Analysis (ROA)

To evaluate the value added as a result of the flexibility created by the option, a “stopping the buy early” option is defined. In our case the “option” is a not an option to purchase *LTB* parts (or additional *LTB* parts), but rather the option to have a *LTB* that is less than the total number of parts needed to satisfy the demand through the *EOS* of the system. Using this definition, exercising the option means making a buy that covers the demand through a time period that ends before the *EOS* date, and letting the option expire means that the *LTB* covers the demand through the *EOS*. We treat the options as a set of European-style options that can only be exercised on a certain date (buy-to date) or allowed to expire. The best “buy-to date” will be the option with the highest value.

The first step in the ROA is to generate demand paths. Monte Carlo sampling of application-specific time-to-failure (*TTF*) distributions for the part is used to generate failure dates for the part and thus the points in time where a spare part is required.⁴ The sequence of demands generated in this way represents a demand path. Having a required demand path, one can find the value of the option at discrete times throughout the life cycle of the system using (2).

Figure 4 illustrates the values generated for a particular demand path. In Fig. 4, the value of the option starts at a minimum where no penalty costs are avoided and all *LTB* costs are avoided (T_{bt1} in Fig. 4). The value reaches a peak where the combination of the penalty cost due to running out of parts avoided and the *LTB* costs for buying and holding parts avoided are maximum (T_{bt2} in Fig. 4). When the buy is made to the *EOS*, all the penalty costs are avoided, but none of the *LTB* costs are avoided (T_{bt3} in Fig. 5).

⁴ Note, in this description we are assuming that there is no manufacturing demand for the part, however, manufacturing demand could be included as well.




Figure 4 - Value from (2) for one demand path as a function of the buy-to time.

Figure 4 shows only one possible value path (corresponding to one possible demand path). In reality, each possible sequence of demands will generate a unique value path, so there will be a large population of possible value paths due to the demand uncertainty. To analyze the population of paths using ROA for the “stopping the buy early” option, we need to define a stopping criteria, (3). The stopping criteria is the value associated with buying the *LTB* all the way to the *EOS*. If the option value is less than the value from buying to the *EOS* date, it lies below the stopping criteria shown in Fig. 5; hence, the option would not be exercised. For example, the option value from (2) at time T_{bt1} is less than that at *EOS* ($T_{bt}=160$ quarters); hence, the option would be allowed to expire. In other words, it is worth more to make an initial buy through the *EOS* at the beginning ($T_{bt}=0$) than having an initial buy up to T_{bt1} and then paying the penalty cost. Alternatively, if the option was exercised at time T_{bt2} , the path has a higher value compared with that at the *EOS*; in this case the option would be exercised at T_{bt2} .

$$\begin{aligned}
 &Value(T_{bt}) - Value(EOS) \geq 0 \rightarrow \text{exercise the option} \rightarrow Value(T_{bt}) = Value(T_{bt}) \\
 &Value(T_{bt}) - Value(EOS) < 0 \rightarrow \text{do not exercise the option} \rightarrow Value(T_{bt}) = Value(EOS)
 \end{aligned}
 \tag{3}$$




Figure 5 - The illustration of the stopping criteria definition for the example case.

For a selected T_{bt} , the defined stopping criteria is applied to each path and the value of either exercising or not exercising the option is determined as,

$$V_i(T_{bt}) = \begin{cases} \text{exercise the option :} & V_i(T_{bt}) \\ \text{do not exercise the option :} & V_i(EOS) \end{cases} \quad (4)$$

where i refers to i^{th} simulated demand path. The final value of exercising the option at a specific buy-to date is obtained by averaging the paths' values, $V(T_{bt})$.

$$V(T_{bt}) = \text{avg}(V_i(T_{bt}) | i = 1:N) \quad (5)$$

where N refers to the total number of paths.

2.3 Demonstration Case

In this section we assume the following application-specific lifetime buy problem:

1. Reliability (for a mean TTF of 7 years):
 - a. The Weibull location parameter = 0
 - b. The Weibull shape parameter = 1.5
 - c. The Weibull scale parameter = 7.7541 years
2. Cost analysis:
 - a. Number of systems to support = 1000
 - b. End of support (EOS)= 40 years (160 quarters)
 - c. Initial buy size = to be determined
 - d. Part purchase price = \$110/part
 - e. Riskless interest rate = 3%/year
 - f. Holding cost = \$38.5/part/year
3. Penalties:
 - a. Under buy (one-time cost) (L) = \$110,000
 - b. Under buy per part penalty (P) = 0
 - c. Overbuy penalty cost = \$55/part

Figure 6 illustrates the result of exercising the “stopping the buy early” option for the example case. The value of exercising this option has its maximum at 93 quarters. Using this time we can determine the required number of parts. Figure 7 shows the plot of 100 simulation demand paths. The distribution of the required buy size at quarters 93 is shown on the left top; the mean buy size at quarter 93 gives 3047 parts as the optimum buy size for the example case.




Figure 6 - The value of exercising the “stopping the buy early” option for the example case.




Figure 7 - The optimum lifetime buy size obtained from exercising the “stopping the buy early”

3. VERIFICATION - STOCHASTIC DISCRETE EVENT SIMULATION (*DES*) SOLUTION

To verify the real options formulation and solution we use a stochastic discrete event simulator. A *DES* is capable of representing the operation of a complex system as a discrete sequence of well-defined events in time (in smaller or larger time scales). Each event occurs at a particular instant in time and marks a change of state in the system.

The *DES* starts with the same demand paths used by the ROA. In the case of *DES*, the demands create a timeline of events that is pushed through a discounted cash flow analysis. The cost analysis includes the initial buy, penalty costs, and holding cost.

Using the example data from Section 2.3, and varying the initial buy size, the result in Fig. 8 is obtained. As it is shown in Fig. 8, by increasing the buy size, the total cost of the system with a fixed *EOS* of 40 years and a fixed WACC of 3%, decreases to a minimum and then increases. The minimum of this curve gives the optimum buy size for such a system. Note that, the optimum buy size is not the mean lifetime demand (e.g., 5445 at 160 quarters in Fig. 7), this is due to the fact that the penalties for overbuy and under buy are not the same; the holding cost is not zero, and the demand is uncertain.





Figure 8 - The total cost versus buy size obtained from exercising DES method for the example

3.1 DES vs ROA

Assuming the same riskless rate and penalties, ROA and a stochastic DES should produce identical results. To verify this, the optimum buy size from the two methods using identical inputs are compared.⁵ To perform the comparison, an identical set of 1000 uncertain demand event paths are considered in both analyses. For a 3% WACC, as illustrated in Fig. 9-a, the DES method gives an optimum buy size range of 2923-3191 with an average of 3052 parts; the optimum buy size for ROA method is obtained from the optimum buy time range; it gives a range of 2922-3190 with an average of 3052 parts. Note, there is a maximum of 1 part discrepancy in this case; this discrepancy is due to time scale differences and rounding numbers. Hence, both methods are consistent in determining the optimum lifetime/bridge buy size.

To further verify this consistency, other WACC values were tested as well. For example, for a 12% WACC, as illustrated in Fig. 9-b, the DES method gives an optimum buy size range of 2502-2817 with an average of 2652 parts; the optimum buy size for ROA method is obtained from the optimum buy time range; it gives a range of 2501-2816 with an average of 2653 parts.



⁵ Note, (2) by itself does not differ from the DES solution. Equation (2) is a “value” calculation, whereas the DES evaluates life-cycle costs. In the case of the ROA, we are interested in maximizing the value from (2), while the DES seeks to minimize the life-cycle cost. Using identical constant WACC values, the two formulations should result in the same optimum buy quantities, since the stochastic DES captures a representative population of possible demand paths.




Figure 9 – The life-cycle buy size distribution for the example case from DES and ROA for WACC

4. DISCUSSION AND CONCLUSIONS

Real options have been used in obsolescence management to assess the value of waiting to invest in new technology, but have never been used to assess lifetime or bridge buys. This report presents a real options approach to value lifetime and bridge buys of parts. Using this valuation approach, the optimum part quantity at which to exercise the “stopping the buy early” option and to perform a lifetime or bridge buy is determined. The method accommodates uncertainties in demand, holding costs, and end of support dates. The optimum buy size from this method has been shown to be consistent with that from stochastic discrete event simulation (*DES*).

The real options analysis methodology is used to address several questions:

- What is the critical initial buy size that drives the option at the money?
- What are the critical penalty costs (under buy, overbuy, and holding cost) that drive the option at the money? How do different sources of uncertainty impact the option’s value?
- What is the sensitivity of the option to uncertainties in the discount rate, *EOS* time, part price, etc.?

A fundamental problem that needs to be addressed is how to set the discount factor (WACC) for the analysis. As demonstrated in Fig. 9, the use of different values of the WACC generate significantly different optimum buy sizes for the same problem. If investors are risk-neutral, then the discount factor is simply the risk-free interest rate; however, it is unreasonable to expect that investors would seek the same rate of return from a project as a risk-free treasury bond. Hence, risk neutrality is not realistic. If investors are risk-averse, they will demand a risk premium on a project. The risk premium can be calculated according to the rate of return on a financial asset that has the same systematic risk as the project. Under the assumption of complete markets (a market in which the complete set of possible bets on future states-of-the-world can be constructed with existing assets), a financial asset with this level of systematic risk exists and so a portfolio or bundle of these financial assets can be created to replicate the risk in the project. In this way, if the project or cash flows were traded, a replicating portfolio can be constructed from the traded project and a risk-free bond. These are the fundamental arguments of real options analysis and would suggest that a risk-free rate should be used when assessing buy sizes.

4.1 The Rochester Electronics Problem

Rochester Electronics suggested a practical case study employing the real options approach to quantify and cost optimization for lifetime buys of parts. Using the Rochester’s historical data, reliability distributions for parts (TTFs, etc.) can be determined. Knowing the TTF for parts, the company needs to

make an initial buy of wafers full of bare die (not finished chips). The purchase of each individual bare die represents the purchase of an option (note, due to the delivery of multiple die on wafer, there is a “stepped” option price here). Wafers can be held until they are needed (or not needed). The strike price in the ROA analysis would be the cost of exercising the option (packaging the bare die into a chip and performing functional testing to enable delivery of the final finished chip). The value of exercising the option is the penalty cost avoidance (penalties due to not having enough bare die) less the cost of exercising the option (packaging and testing the chips). Finally, the options would expire if the chip is not needed (in this case the strike price is not paid and the bare die is not packaged or tested).

The key question in the Rochester problem is how many die should be optioned (i.e., how many wafers should be purchased in this case). Clearly, each possible demand path for this problem has the flexibility to decide whether to exercise each (die purchase) option or not. Conventional approximations using stochastic discrete event simulation cannot accurately model the Rochester problem because they cannot model individual path flexibility. ROA simulations and possibly lattices (e.g., a binomial method) can be used here. Using a ROA simulation technique, one can generate a cone of uncertainty, the project payoff at different times throughout the lifetime of the system, and finally obtain the optimum buy size.

There are several complicating factors in the Rochester problem. First, in order to treat this problem with a lattice we will have to assume that the wafer holding costs are negligible (or absorbed into a uniformly applied overhead rate) – Rochester has indicated that this is a good assumption. Second, the options can only be purchased in bundles (or batches), not individually, because you are paying for a wafer full of die. A third issue is that the yield of the bare die on the wafer is not 100%, but you will not be able to fully test the die until they are packaged, therefore, the strike price (cost of packaging and functional testing) has to include the cost of packaging and testing some unusable die, i.e., the yielded cost (Becker and Sandborn 2001) of packaging and testing has to be used rather than the actual cost. Note that a fixed cost of the manufacturing and testing facilities (independent of the initial buy size of wafers) has to be added to the strike price in the ROA analysis. Using the ROA simulation technique, the uncertainties in the required number of options based on the variable demand throughout the lifetime of the system and also the uncertainties in the possible yield of the bare die on the wafers are addressed in a stochastic approach. Hence, the optimum initial buy size of wafers or buying options is determined.

While the holding cost, buying options in bundles or individually, delayed determination of the yield of the bare die, and the one-time manufacturing and testing facilities cost appear in the Rochester and similar problems, there are more specific aspects that should be addressed in the case-specific Rochester problem. First, there are a limited number of companies that produce required packaging materials (i.e., package bodies - four companies in this case). So the problem is not limited to options for bare die, it also has to consider options for package bodies – this could be a more limiting factor. Second, the type of interconnection on the chips (e.g., wirebond or flip chip) and the technology transition should be considered in the valuation of options. Different types of chips may have different geometries that would result in a different fabrication process/cost. The technology drives different testing facilities and may require hiring personnel. These all result in increases and complexity in the strike price.

5. REFERENCES

- Becker, D., & Sandborn, P. (2001). On the use of yielded cost in modeling electronic assembly processes. *IEEE Trans. on Electronics Packaging Manufacturing*, 24(3) 195-202.
- Burnetas, A. N., & Ritchken, P. (2000). Option contracts in supply chains. Technical Memorandum Number 728, Department of Operations, Weatherhead School of Management, Case Western Reserve University.
- Edgeworth, F. Y. (1888). The mathematical theory of banking. *Journal of Royal Statistical Society*, 51, 113-127.
- Feng, D., Singh, P., & Sandborn, P. (2007). Optimizing lifetime buys to minimize lifecycle cost. *Proceedings of the 2007 Aging Aircraft Conference*, Palm Springs, CA, USA.

- Josias, C. L. (2009). Hedging Future Uncertainty: A Framework for Obsolescence Prediction, Proactive Mitigation and Management (Doctoral Dissertation). University of Massachusetts, Amherst, USA.
- Sandborn, P. (2008). Trapped on technology's trailing edge. *IEEE Spectrum*, 45(4), 42-58.
- Sandborn, P., & Myers, J. (2008). Designing engineering systems for sustainment. *Handbook of Performability Engineering* (K.B. Misra, Ed.). London, Springer.
- Sandborn, P. (2013). *Cost Analysis of Electronic Systems*, World Scientific Publishing Co, 307-328.
- Stogdill, R. (1999). Dealing with obsolete parts. *IEEE Design & Test of Computers*, 16(2), 17-25.
- Teunter, R. H., & Fortuin, L. (1999). End-of-life service. *International Journal of Production Economics*. 59, 487-497.
- Teunter, R. H., & Haneveld, W. K. (1998). The "final order" problem. *European Journal of Operational Research*, 107(1), 35-44.
- Venesmaa, J., Hietaniemi, Tiainen, O., Roine, T., & Niinistö, A. (2008). Real options and demand uncertainties in supply chain management. Helsinki University of Technology, Finland.

